DOCUMENT RESUME

ED 307 104 SE 050 523

AUTHOR Cocking, Rodney R.; Mestre, Jose P.
TITLE Cognitive Science.
PUB DATE 89
NOTE 23p.
PUB TYPE Information Analyses (070)

EDRS PRICE MF01/PC01 Plus Postage.
DESCRIPTORS Artificial Intelligence; *Cognitive Processes; Cognitive Psychology; *Learning Theories; *Problem Solving; *Schemata (Cognition)
IDENTIFIERS Chunking; *Cognitive Sciences; *Expert Novice Problem Solving

ABSTRACT

The focus of this paper is on cognitive science as a model for understanding the application of human skills toward effective problem-solving. Sections include: (1) "Introduction" (discussing information processing framework, expert-novice distinctions, schema theory, and learning process); (2) "Application: The Expert-Novice Paradigm as a Means of Studying Problem-Solving" (describing chunking, hierarchical memory networks, expert-novice differences in problem solving style, promoting expertise, and artificial intelligence tutors); and (3) "New Directions and Future Trends" (a brief look at promising developments including sensory processing approaches, developmental sensory integration, and robotics). This paper contains a list of 71 references. (YP)

****
* Reproductions supplied by EDRS are the best that can be made *
* from the original document. *
****

ED307104

SRRI #211

To appear in Handbook of Knowledge Engineering (1989). Washington,DC: Systemsware.

36. COGNITIVE SCIENCE

"PERMISSION TO REPRODUCE THIS MATERIAL HAS BEEN GRANTED BY
Jack Lockhead
TO THE EDUCATIONAL RESOURCES INFORMATION CENTER (ERIC)."

Rodney R. Cocking, Ph.D.
National Institute of Mental Health
and
Jose P. Mestre, Ph.D.
University of Massachusetts

U.S. DEPARTMENT OF EDUCATION
Office of Educational Research and Improvement
EDUCATIONAL RESOURCES INFORMATION CENTER (ERIC)
☒ This document has been reproduced as received from the person or organization originating it.
☐ Minor changes have been made to improve reproduction quality.
• Points of view or opinions stated in this document do not necessarily represent official OERI position or policy.

Introduction

The relevance of research in Cognitive Science to the focus of this handbook, Knowledge Engineering, is in its contribution to the understanding of problem solving processes. Cognitive Science has links to Artificial Intelligence, to cognitive psychology, to information processing, to language-based information systems, and a variety of other areas. For this discussion, however, the focus will be on Cognitive Science as a model for understanding the application of human skills toward efficient, effective problem-solving. Knowledge engineering questions lie within a framework of information processing and of how a comprehensive analysis of critical skills can assist in moving novice performance to expert performance in as efficient a manner as possible.

The Cognitive Science model has been applied most broadly at the variable level for analyzing the scope of a problem and for specifying the performance skills that relate to each variable. For example, when a research team wanted to analyze the relationship between school performance on mathematics achievement and the students' language skills, the questions could have been addressed by considering individual topics that relate to math and language performance. Application of a cognitive science model, by contrast, began by addressing the problem state (what was known) and the goal state (what was to be learned) and then framing the problem broadly in terms of relevant variables (Cocking & Chipman, 1988).

The beginning point in this analysis was to take a systematic look at the relationships between math achievement and language status variables. This approach required an examination of the relevant variables that relate to these children's conceptual, developmental, and linguistic status for receiving and utilizing classroom instruction. The aspects of the problem were schematized along the lines of Input to the children and Output (i.e. child performance). On the Input side are Cognitive Ability Patterns (including math learning, language skills, reading); Educational Opportunity (including time on math tasks, quality of instruction, receptive language, parental assistance, parental education); and Motivation to Engage (including cultural values, parental influence, expectations for reward, motiviational nature of instruction interactions, equitable treatment). On

SE 050 533

BEST COPY AVAILABLE

the Output side are Measurement Issues (including sensitivity to developmental status, culture fairness), Lanaguage of Test (including instructions for what to do), and Performance Variation (including types of math problems - word versus computational problems, and math versus other cognitive skills of performance).

Input/Output variables, however, frame the problem only in the most global terms. Specific skills are associated with the array of variables and it is at this next level of analysis - at the level of information processing skills - that the Cognitive Science model builds upon research from cognitive and developmental psychology and where implications for Knowledge Engineering emerge. A brief overview of the variety of cognitive skills that are important in the Cognitive Science model will be laid out next, and then specific focus will be directed toward problem-solving. Problem-solving is by no means the only aspect of the cognitive science model that applies to Knowledge Engineering; however, this discussion will be limited to the problem-solving issues, with occasional contrasts brought in from other related areas, such as learning.

An Information Procesing Framework

The starting point for the information processing framework is to ask, What are the basic behavioral proceses that enable humans to make sense out of environmental information? This means specifying how humans attend to information, select critical information, and interpret their environments. In short, how do people respond in an orderly way to their environments, as opposed to dealing with a "scrambled" world? Processes of attention, perception, emotion, and language are basic mechanisms for filtering the environment. What are the associated learned skills that utilize these basic mechanisms?

The Information Processing framework is useful for identifying the critical behavioral processes. The framework is sufficiently broad to allow casting these questions from perspectives of the social environment, the emotional environment, and the intellectual/cognitive environment. For example, major concerns within behavioral science include how people learn the information that is essential to their adaptation and mental growth, how they store or remember experiences, and how their performance is improved or becomes more adaptive. The information that is learned, remembered, or used can be social information, information about how they feel (their emotions), or cognitive skills. The basic processes of language, perception, comprehension/interpertation, categorical grouping, attention, etc, apply across domains of information. In a listing, then, the following questions group into 4 major classes that imply different skills and operations: (1) How is information learned? (acquired); (2) How is information stored? (represented or encoded); (3) How is information remembered? (retrieved or decoded); and (4) How is information used? (applied).

The framework suggests looking at components of behavior, such as the learning issues, the encoding or storage issues, the memory issues, and perhaps most relevant to our discussion of problem-solving, the performance

issues. This framework links performance to the other issues, such as to learning and to memory so that one can begin to target sources of performance enhancement or deficits at the learning (acquisition) or the memory (retrieval) stages of processing information. The same questions apply regardless of the domain of knowledge, since the underlying details we want to know are about the behavioral processes of learning, namely how information is stored, how information is remembered, and how performance is improved with learning experiences. In addition, the basic behavioral processes are further layered onto these questions by addressing concerns such as how attending (attention processes) selectively influences understanding an event. Thus, processes of selective attention, awareness, inferences, and interpretation apply to any event, whether the event is social, affective, or purely cognitive in nature.

In summary, the framework breaks down into learning, memory (two aspects being representation and retrieval), and application or usage of skills and information. Problem-solving is largely the last component, though it has obvious links to learning and memory for achieving expert status of efficient and effective performance. Next we will consider the relevant aspects that lead to expert performance and follow with a discussion of some additional critical aspects of performance that set experts apart from novice problem-solvers.

Performance Differences

The most rudimentary, initial phase of analyzing a problem is to characterize the subject matter -- the domain. While there are general intellectual skills that are generic to all or most intelligent behavior, identifying the domain of problem-solving (mathematics, geometry, measuring ingredients before baking a cake) is to distinguish the relevant information pertaining to the task at hand. Second, it is critical to impose some organizational structure on the relevant skills. Thus, the beginning point is to acknowledge domain specificity, and further to acknowledge that there are identifiable, qualitative differences in ways for performing the task that set experts apart from unskilled or semi-skilled problem-solvers.

The Expert-Novice distinction can be apparent in many forms: in how easily or laboriously learning occurs; in differences in how information is represented; in storage differences; in terms of access or retrieval efficiency; and in terms of skilled application or vision for potential application of domain-specific knowledge. Experts differ in terms of their information levels, as would be expected, but additionally they appear to employ more sophisticated and efficient learning strategies, they categorize the problem types and analytic strategies for solving a problem differently from novices; they remember problems in different ways; and they go about problem solutions differently. For example, Novices tend to work backward from goal states in their attempts to construct solutions to problems. As a consequence, they may perform adequately by solving a problem, but their performance is oftern so problem-specific that it may have little or no generality (transferability to new situations). Experts, by contrast, work forward, starting from general principles and moving toward more specific procedures for applying the principles; in doing so, they thoroughly explore

the problem space. In addition to considering a wider array of variables that contribute to the problem, Experts end up learning about types of related problems in a domain and hence, experts are different from novices in both learning and problem-solving skills.

It is important to think about the Expert-Novice distinctions in order to address the differences between novice performance and expert performance. Although the Expert-Novice distinctions discussed above appear to be too general to be of any practical use, we will re-visit them below in two different contexts and with increasing specificity. Since the task of Knowledge Engineering is to move novice performance toward skilled, expert-like performance, it is necessary to ask, first, what is it that experts do and then to ask how one can move a novice toward the desired performance level. This leads us to two important considerations in Cognitive Science: Schema theory and general issues of Learning.

Schema Theory

Information has to be organized in some way in order to be meaningful or useful. Cognitive psychologists have developed models of how knowledge is represented in people's minds (e.g., Bower, 1975; Rumelhart & Norman, 1975; Schank, 1975; Schank & Abelson, 1977). These models help define the kinds of knowledge people have, how knowledge is acquired, how people retrieve knowledge, and how information is used (Dehn & Schank, 1982). One theoretical model for how humans learn about organizational features of events in their worlds is schema theory.

Schank & Abelson (1977) asked a simple question to guide their model: "What do we know about typical life events that we use for making inferences and predictions?" They devised a descriptive model of what people know about typical events so that inferences can be drawn. The Schank & Abelson model, termed schema theory, specifies sequences of actions that are linked temporally and causally. The key element in the model is the script, which is "a basic level of knowledge representation in a hierarchy of representations that reaches upward through plans to goals and themes" (Nelson, 1982, p. 101). A scheme describes a present-state framework into which actions are organized, in effect a structure into which new information can be incorporated or accommodated (Sigel & Cocking, 1977; Cocking, 1983). The scripts or schemes are well-specified and are concrete, in contrast to abstract levels of goals and themes (Nelson, 1982).

Experimental studies have used this model for studying memory and results indicate similar recall organization for children (Nelson, 1978) and adults (Bower, Black & Turner, 1979). These memory studies illustrate similarities in event elements, event structure, sequences, references to implied but unstated related elements, common inferences, and general rather than specific episodic themes. Scripts, therefore, are general in form, temporally organized, consistent over time, and socially accurate (Nelson, 1982, p. 103). Research in recent years has expanded the model to include how scripts, as generic organizational frames, are acquired and applied. Research on application of schema theory-related topics includes concept

development, conceptual thinking, classification skills, organizational strategies, etc. (Gerhard, 1975; Cocking, 1983; Friedman & Cocking, 1986). In all of these manifestations, the underlying argument is that organization of information is a critical device for knowledge represention and storage. Experts utilize broader "chunks" as they encode information and these chunks are semantically organized (that is, they are meaningful units). Experts, further, can identify more subcomponents to the broader group; that is, the storage system is more efficient in that it subsumes related informational units. In essence, efficiency for storage and retrieval means that experts are capable of thinking and remembering in larger conceptual, categorical classes into which lesser information can be incorporated. The net effect of this approach, and that which distinguishes experts from novices, is, in a word, efficiency.

A concrete example illustrates the application of schema theory in teaching organizational and communication skills. Gerhard (1975) uses paragraph writing as a way of teaching categorical thinking strategies. One first considers the range of elements to be included in presenting some information on a topic. The next task is to decide which one of the list of, say 6 items, is to be the organizing theme and to write a topic sentence. The net effect of writing a sentence about each of the 6 elements after deciding which is to be the organizing item leads to a topically organized sequence of related items. In essence, schema theory posits that this is how efficient storage-encoding and memory-decoding operate: only the organizational category needs to be encoded or retrieved which results in cognitive economy of informational structure. This, in effect, is also the operational definition of a schema -- a superordinate organizational framework.

Learning

The discussion implies that schemas are critical components of problem-solving. Equally obvious, then, is the importance of schema acquisition to Knowledge Engineering, since the goal is to move people in the direction of becoming effective problem-solvers.

One possible mechanism for influencing learning and cognition is instruction. Vygotsky (1978) argues that children learn to solve problems through opportunities to solve them with more expert individuals. These experts structure the problems to be only slightly more difficult than problems the child can solve on his or her own, and they direct the child's problem-solving so as to allow the child to function at the upper ability limits. Some (Feuerstein, Rand, Hoffman, & Miller, 1980) posit that parents mediate experiences for children by giving new experiences structure and by responding to aspects of the environment so that there is a more systematic, planful and logical structure. Research indicates that mothers help children solve problems, classify objects, and prepare and check memories for what is being learned.

But not all learning is influenced by experts who show or tell novices

how to reason, plan, or solve important questions, however. Considerable human learning occurs without instruction, through sensory systems 'wired' into the human neural network for picking up critical environmental information (J. J. Gibson, 1966). E.J. Gibson (1982) writes that perceptual-based learning is "initiated through exploration motivated both intrinsically and extrinsically." Piaget's work focuses on many aspects of learning that occurs with relatively little intrusion from expert instruction or guidance (Sigel & Cocking, 1977). Thus, cognitive psychologists who extoll the virtues of exploration and learning by discovery are making a distinction between knowledge that is directed toward a specific goal (knowing the date of an historical event) and non-specific knowledge goals of general characteristics of a domain (determining if an ice cube melts faster sitting on a saucer or in a dish of room-temperature water). This distinction is an important one in effective problem-solving, as will be pointed out in a subsequent section.

Thus, most people consider it self-evident that instruction is necessary or that it is at least the most efficient manner of transmitting values and for conveying the symbol systems of a culture (Gardner, 1986; Gardner, 1984). Reading, writing, and mathematics most probably cannot be picked up without instruction. Specific systems of reasoning and problem-solving such as those represented by the scientific method probably cannot be acquirted without instruction and tutoring. Evidence from cross-cultural research and from training studies clearly demonstrates that instruction facilitates learning and cognitive functioning. This is true for a variety of learning tasks, including perceptual, memory, and logical reasoning tasks, free-recall, classification, recognition memory, and so forth (see Rogoff, 1981 for a detailed review). While much of the debate surrounding whether schooling actually alters cognitive development and results in improved cognitive functioning, the fact is that instruction has been shown to change performance on specific cognitive tasks. Specifically, instruction has been shown to lead to improved cognitive performance on tasks such as physics and mathematics problem-solving (Larkin, Heller & Greeno, 1980); writing (Bereiter & Scardamalia, 1978; Gerhard, 1975); reading (Brown, Campione & Day, 1981); and cognitive skills such as thinking, problem-solving, and reasoning (Meichenbaum, 1977; Nickerson, Perkins & Smith, 1980).

It is commonly assumed that practice on a large number of typical problems is the optimal method of acquiring problem-solving skills. That is, rote learning-by-doing methods of education are assumed to be the best training exercises. Sweller (1988) and others question this assumption, given what Cognitive Science research models have begun to indicate about how domain-specific knowledge affects problem-solving. There is reason to believe that practice skills are useful for certain learners (e.g., children who have a limited skills repertoire and a limited arsenal of past experiences from which to draw), but the Cognitive Science literature for adult learners clearly challenges the myths of routine practice exercises as an efficient means of promoting expertise.

Owen & Sweller (1985) showed that nonspecific goal-oriented problem-

solving could be contrasted with conventional means-ends problem-solving strategies to illustrate critical, meaningful differences between the two approaches. The former approach is termed "forward working," whereas means-ends problem-solving is termed "backwards working." Forward working is an expert system that is schema-driven. In this approach, the problem state is analyzed and the necessary operations as well as options are specified for moving toward a problem-solution of a goal state. This type of analysis sets up a contrast between present conditions with what is needed to reach a specified desired set of circumstances. A means-ends approach, by contrast, tends to invoke all of the steps of achieving a goal without regard to redundancy or irrelevancy of certain steps. A forward-moving approach also focuses upon possible alternatives through exploring the problem and discovering features that make one problem type different from another. Considerations of this sort are generally avoided if one's eyes are fixed only on an end goal state of a means-ends strategy. Sweller and his colleagues (Sweller, et al, 1983) replaced specific goal-oriented directions in a problem set with non-specific goals (e.g., "calculate as many variables as you can" versus "calculate a race car's acceleration"). The effect proved beneficial in schema acquisition, that is in learning. While the practice may be questionable for performance (that is, it is not as fast or as direct), this alternative approach clearly taught the learners more about the task than the traditional set of instructions. The net effect was faster acquisition of expert-like schemas and schema-driven approaches to problem solving.

Sweller accounts for this learning efficiency in terms of the reduced memory load required for forward working problem-solving, as compared to what is required in means-ends approaches. From this and other studies conducted by the group, Sweller concluded that "Problem solvers [who organize] a problem according to means-ends principles, suffer from a cognitive overload which leaves little time for other aspects of the task. The overload can be manifested by an increase in the number of ...errors" (Owen & Sweller, 1985; Sweller, 1988, p. 276). Sweller also points out that trying to learn problem-solving strategies (acquiring schemas) at the same time as solving problems via means-ends strategies is akin to doing two things at once. Solving the problem is the primary task, while trying to figure out what may be useful to know in the future is a second task. The question, then, is whether there a dual task when one is "learning by doing?"

The test of the "dual task theory" was to look at both phases of performance. If the primary phase of means-ends strategies places a heavy cognitive processing load, the net effect should be lowered support or interference effects for the secondary phase. By contrast, if experts use cognitive classification strategies to remember and access meaning for similar types of problems, then one might expect facilitation across the two phases. Experimental results indicated no differences in the total time it took to solve the problems, but that there was a heavier cognitive processing load in the conventional means-ends approach--that is, the dual task interferred with the secondary task unless a more expert-like schema-driven strategy was employed. The conclusion was that there was more learning during "doing" (that is, during performance) when

the doing was schema-driven.

This raises the question, Why is there interference? The answer lies in the differences between learning schemas and means-ends analyses. Two critical mechanisms underlie learning and problem-solving: (1) attention and (2) processing capacity. In the Cognitive Science model, acquisition of information and application of information are distinct processes; hence, learning and problem-solving are seen as largely unrelated processes. For means-ends analyses, a problem solver attends to differences between a current problem state and a goal state. Previous states, in this framework are relevant only for preventing repetitions or retracing steps. By contrast, in schema acquisition a problem solver needs to be able to recognize and classify a problem as belonging to a class of problems. Within this framework, previous problem solutions represent a problem state for that particular problem and as such a represented solution may be an important component of problem-solving. But the goal state is not the only critical feature for understanding the problem (as it is in a means-ends representation). What is important in schema theory is the role of representations in memory for purposes of acquisition. That is, there is a structure into which new information can be incorporated. Piaget terms the acquisition of new information "assimilation," while the molding of that new information into an existing structure is an "accommodative" process. By this account, then, variables of structure that relate to perception and perceptual pick-up are likely to be most associated with schema acquisition, while variables related to meaning are most relevant for encoding and retrieval in performance tasks.

Memory is also important to schema acquisition in that categories are memory-related, but this is not to say that learning categories of information is dependent upon memory. This distinction, while subtle, is a radical one that is revolutionizing instruction. For example, math teachers can present a problem set which capitalizes upon classroom examples which a student has to recall, recognize, and remember for their unique or relevant features. Such a routine is preferred to repetitive practice sets because it draws upon memory for matching past problem types with current demands. A math problem which requires youngsters to make some measurements and to convert those measurements into fractions prior to setting up the problem for solution means that the student assesses problem-state and goal-state and identifies relevant variables without recourse to memory skills of matching present end-goal states with remembered goal-states. In the Cognitive Science model, therefore, mnemonic skills (memory) are relatively more important to schema acquisition than to problem-solving. Memory, in fact, may be an important source of interference in performance or schema-application when the wrong problem prototype is pulled from memory. This conclusion is certainly counter-intuitive to those who believe that performance is based upon matching a problem-solving strategy to the problem type as classified on surface-level problem features.

An example of this "matching to memory" strategy and the errors it can lead to has been shown in math word-problems (Clement, 1982; Mestre, 1988). The structure of English grammar leads one through a left-to-right

processing strategy that corresponds to the left-right reading scheme. This surface-level feature, encoded as a basic memory unit, often leads to errors for solving a certain class of algebra problems, since English syntax often transposes subject and object, or in algebra problems, the x and y variables. Mestre presents the following example that evokes the tendency for using the left-to-right reading strategy to parse an algebraic statement into an algebraic equation: "Write an equation using the variables S and P to represent the statement 'there are 6 times as many students as professors at this university.'" The common error, termed the variable-reversal error, consists of the answer 6S=P, even though the respondents acknowledged that there were, in fact, more students than professors. The error stemmed from focusing on the problem's surface features coupled with the highly practiced habit of parsing a sentence using a sequential left-to-right strategy. If students had focused on the problem's deep structure, this error could have been avoided. Memory, as invoked in means-end strategies by novices, illustrates how selective attention to surface level details can lead to misclassifying both the problem type and relevant strategies.

Another source of interference during learning is from the demands placed on the cognitive system -- or the cognitive load. Means-ends analyses used by novices also have a "cognitive load" associated with them that potentially interferes with learning. Means-ends solutions may be "efficient" in the sense that the strategy generally leads to few dead ends. However, Sweller contends that there is a price for this efficiency: the strategy requires the problem solver simultaneously to consider the problem state, the goal state, the subgoals, and all the problem-solving operators. The net effect of this coordination is a heavy toll on a limited procesing capacity. Sweller believes that all of this effort leaves little cognitive capacity for attention to schema acquisition (learning). So much effort is devoted to problem-solving that little cognitive capacity is allocated to learning and exploring the problem space.

Evidence that there is such a thing as cognitive load has been obtained by looking at a number of criterial variables relating to working memory. Evidence for increased, excessive, and unnecessary cognitive load imposed by means-ends approaches in problem-solving has come from analyzing the kinds of strategies employed; the categories of viable, usable solutions; the speed of solution; errors in subsequent problems; and modeling techniques (Sweller, 1988). It should be pointed out, however, that some of these criteria fit performance better than learning objectives. Is speed of problem solution critical? It depends. In planning-related tasks (Friedman, Scholnick & Cocking, 1987) the criterion of efficiency in carrying out performance was often preferred to a criterion of speed when accuracy or appropriateness was not altered. Error types, another criterion, may reflect developmental status of the task performer. In an example such as solving a language-related problem, some error types are "more sophisticated" and represent advanced developmental stages of language growth than earlier errors, though both, on an absolute scale of right/wrong, are classified as errors. Cognitive load and processing capacity, therefore, have to be considered relative to standards for specific problems and problem-solvers. That is, not all problems are of equal difficulty and expert systems, while specifying no age-relatedness,

generally imply adult status.

Application: The Expert-Novice Paradigm As A Means of Studying Problem-Solving

In this section we elaborate on some of the general concepts discussed thus far. Again, the focus is on comparing and contrasting the behavior of Experts and Novices in various cognitive tasks. The examples provided below should illustrate how Cognitive Science arrives at the models of knowledge acquisition, knowledge representation and retrieval, and knowledge usage discussed above.

The goal in studying the end points of the expertise dimension is to gain insights on the salient features of expertise and novicity, and thereby gain insights on efficient instructional methods for moving novices toward expertise. Although indications are that expertise is very context-dependent (Brown, Collins & Duguid, 1989; Perkins & Salomon, 1989) and generally does not transfer from one domain to the next (e.g., an expert mathematician will not be able to use his or her mathematical expertise in the domain of chess and vice versa), results from expert-novice studies are generalizable across domains. That is, there are many commonalities in the way that experts from different domains acquire, store and use domain-related knowledge to solve problems.

Historical Beginnings -- Chunking

Findings from expert-novice studies offer suggestions for the design of efficient instructional approaches in promoting expertise. Some of the first studies of expertise were conducted in the domain of chess (Chase & Simon, 1973; de Groot, 1965; Newell & Simon, 1972). The task that separated strong from weak players was a memory recall task where players were shown a chess board configuration for a very short period of time and asked subsequently to reproduce as much of the board configuration as they could. Experts were able to reproduce the position of the majority of the chess pieces on the board, whereas weaker players could not match the experts' recall ability. Memorization ability had to be discarded as an explanation because strong and weak players were equally poor at recalling board configurations made by randomly arranging chess pieces.

These findings have been reproduced in the domains of electronics (Egan & Schwartz, 1979) and computer programming (Ehrlich & Soloway, 1984). For example, expert electronic technicians are capable of reproducing large portions of complicated circuit diagrams after brief exposures, whereas less experienced technicians could not. Similarly expert programmers could recall large sections of programs after brief exposures whereas novice programmers could not. As in the chess experiments, skilled electronic technicians did not hold the same advantage over novices in the recall of circuits composed of randomly arranged symbols, and expert programmers were as poor as novices in recalling "nonsense" computer programs composed of a series of randomly arranged programming statements.

These differences were explained in terms of experts' ability to group together clusters of information according to some underlying principle or pattern. In chess, the experts grouped together clusters of chess pieces according to some underlying strategy or goal of the game, and thus recalled the board configuration by first recalling the strategic clusters and then the individual pieces within each cluster. In electronics, the experts grouped together clusters of individual components (e.g. resistors, capacitors, diodes, etc.) into functional unit clusters (e.g., amplifiers, rectifiers, etc.). The mechanism by which items are grouped by some underlying goal or principle is called chunking.

Chunking has also been observed in problem solving tasks among expert physicists (Larkin, 1979). Expert physicists engaged in solving classical mechanics problems could generate clusters of relevant equations in spurts, suggesting that these clusters were accessed in memory via some underlying principle or concept. In contrast novices generated equations individually with time gaps separating each equation generated.

From Chunking to Hierarchical Memory Networks

The chunking mechanism observed in experts is a precursor to the schema theory described earlier which is so useful in describing how experts' knowledge is stored in memory. The experts' domain knowledge in memory can be thought of as a hierarchical network where there is a pecking order of importance associated to where a piece of information is stored in the hierarchy. At the top of the hierarchy are a small number of "umbrella" concepts to which are attached relevant ancillary concepts, facts, and procedures for applying related knowledge in problem solving situations; the umbrella concepts and their associated declarative and procedural knowledge can also be described in terms of schemas (or schemata). Unlike experts, the novice's memory store is more amorphous in structure.

That experts have a conceptual hierarchy is also manifested in experiments of problem categorization. One commonly used paradigm for problem categorization experiments is to give the subject a stack of cards each containing a typed problem, and ask the subject to place the cards into piles according to similarity of solution; that is, those problems that would be solved similarly should end up in the same pile. Findings from card sorting experiments reveal that experts cue on the underlying concept or principle that needs to be applied to solve the problem as the categorization criterion (Chi, Feltovich & Glaser, 1981; Hardiman, Dufresne & Mestre, in press; Schoenfeld & Herrmann, 1982); this is referred to as cuing on problems' deep structure. For example, expert physicists will place problems requiring the application of Newton's Second Law into one pile, problems requiring the Work Energy Theorem in another pile, and so on. In contrast, novices appear to cue on problems' superficial similarity in deciding solution similarity. For example, novice physics students tend to make problem piles in which the problems share some common object, terminology or other superficial attribute (e.g. problems having to do with inclined planes are placed in the same pile, those having to do with friction are place in another pile, those having to do with rolling objects

are placed in a third pile, and so on). For obvious reasons, this is referred to as cuing on problems' surface features.

Expert-Novice Differences in Problem Solving Style

Glaser and his collaborators (Chi & Glaser, 1981; Glaser, 1984, 1989; Rabinowitz & Glaser, 1985) discuss a profile of expertise that incorporates many of the constructs discussed already. Among the salient features of expertise are a large, highly specialized knowledge base, the rapid perception of meaningful patterns and fast access to relevant knowledge from memory appropriate for the recognized patterns, a rich arsenal of procedures for implementing principles in a forward fashion from givens to goals, and a self-monitoring mechanism by which experts regulate/evaluate the validity of problem solving moves.

In quantitative domains such as mathematics and physics, many of the Expert-Novice distinctions in problem-solving style become readily apparent. For example, novices clearly display their tendency for using the backward-working means-end analysis when solving problems (Larkin, McDermott, Simon & Simon, 1980a, 1980b; Simon & Simon,1978). Unlike novices, experts' problem solving style is more forward-working and principle-based. Experts appear to begin the process of constructing a problem's solution by performing a qualitative analysis of the problem in terms of principles and heuristics that they may wish to apply (i.e., they use the problem's deep structure in performing the qualitative analysis). The result of this analysis can be thought of as a high level strategic map that allows the expert to efficienty move forward from the problem's givens and from the selected principle(s) and heuristics toward the goals. That ability to perform qualitative analyses is an expert trait was illustrated in a study where experts and novices were asked to articulate the approach they would use to solve problems (Chi, et al., 1981). It was found that experts were very eloquent in stating the principle that they would apply and what procedures they would use to instantiate the principle, whereas novices did not discuss a strategy for solution; rather, novices jumped into the solution itself, stating equations they would use in solving the problem without discussions of general principles or procedures.

Promoting Expertise

At this point the attentive reader might reason that a quick and easy method for moving novices toward expertise is to make them aware of experts' characteristics and the powerful methods that they use to solve problems. This approach is naive. Simply telling novices what the expert's characteristics are does not mean they will be able to employ them or emulate them. The arsenal of procedures possessed by experts to solve problems is tied to a rich knowledge base. Thus, teaching novices generalized expert-like heuristics, even if they understand them and are eager to apply them, is insufficient if they do not know how and when to apply them within the context of the domain. Powerful problem solving techniques must be accompanied by the knowledge base within which to apply the techniques (Schoenfeld, 1985). Recent research on instructional approaches based on cognitive research findings indicate that there are more

efficient methods for helping novices reach expert status.

One study by Eylon and Reif (1984) investigated the influence of the form in which a physics argument was presented. One group of subjects received the argument in a hierarchical format (i.e., high level concepts, procedures and goals were separated from information deriving from the high level information), while another group received the same argument in a linear, nor-hierarchical format. The group receiving the argument in hierarchical form performed significantly better on both recall and problem solving tasks. These results suggest that the organization of a presentation can be as important as its content in terms of people's ability to assimilate it in meaningful chunks and use it in problem solving settings.

In another study Heller and Reif (1984) trained novices to generate qualitative analyses of physics problems involving Newton's Second Law before they were allowed to solve the problems. Subjects were trained to perform a detailed analysis of a problem before attempting a solution, to determine what relevant information should go into the analysis of a problem, and to decide what procedures can be used in carrying out the solution plan. This training resulted in substantial improvement in ability to construct problem solutions. These researchers attributed the success of the treatment to the explicit teaching of qualitative analyses that precede experts' problem solving, and accurately point out that qualitative analyses are seldom taught in physics courses.

Finally, more recent studies (Mestre, Dufresne, Gerace & Hardiman , 1988; Touger, Dufresne, Gerace & Mestre, 1987) investigated the possibility of promoting expertise with a treatment encompassing all of elementary classical mechanics. The treatment used in this study consisted of constraining novices who had performed reasonably well in a mechanics course to follow a hierarchical, top-down analysis of physics problems. This expert-like analysis began by asking the subject to select a fundamental principle that could be applied to solve a problem under consideration. After selecting a principle, the subject had to specify the principle further (e.g. select ancillary principles and concepts), and instantiate the principle through some appropriate procedure. No quantitative information (i.e. equations) appeared throughout the analysis until the analysis was completed; at this point, the subject was shown the principle, and procedure used to instantiate it, in equation form. This equation, or set of equations, could then be used to generate a solution to the problem. In order to streamline the analysis, the hierarchical approach was implemented in a menu-driven, computer-based environment. Subjects showed significant improvements when compared to control subjects in ability to categorize problems according to deep structure and ability to draw on principles in performing a qualitative analyses of problems. What is interesting about this study is that subjects were neither trained to use the approach nor provided with feedback to help them ascertain whether or not they were actually performing the analysis correctly. Subjects were simply exposed to the top-down, principle-based approach.

It therefore appears that exposing novices to approaches based on

applying principles and procedures in a forward manner helps them appreciate the role of principles in problem solving. A more recent extension of this work indicates that those who stand to reap the most benefit from this approach are novices who initially show medium to low proclivity toward cuing on deep structure in problem categorization (Dufresne, Gerace, Hardiman & Mestre, in preparation). This is encouraging since it implies that those who improve most from exposure to expert-like approaches to solving problems are novices who exhibit the least expert-like behavior.

Artificial Intelligent Tutors

An entirely different approach to promoting expertise combines the power of technology with the advances in cognitive science: Artificial Intelligent Tutors. An AI tutor is a computer-based system that "reasons" about the learner and tailors instruction to maximize learning. As such, an AI tutor must model four separate entities: the domain knowledge, the communication environment (control system, screen design, menus, windows, etc), the cognitive processes of the student, and the tutoring strategies. The domain knowledge is modeled with the help of domain experts. To build an AI tutor that teaches problem solving in physics, the design team would include expert physicists whose job is to articulate all relevant knowledge needed to teach the desired skills. Once this knowledge is articulated, it must be coded and represented so that the computer can use it to reason about the domain. For example, the tutor must be able to decide whether or not a student's course of action is appropriate for solving a particular problem in order to decide whether to leave the student alone, or to interrupt with some intervention strategy.

Modeling the communication environment often involves a team of computer scientists, human factors engineers and cognitive psychologists. In modeling the environment, the team must decide how the system will "hang together." Questions that must be answered before the communication environment is modeled include: What actions will the student be allowed to make?, What actions will the computer make in communicating with the student?, How will the learning environemnt look (i.e., will it be a multi-window menu-driven environment, will it include graphics and simulations, will it let the student explore the problem space, etc), How will the domain knowledge, the communication environment, the cognitive processes of the student, and the tutoring strategies be linked? (i.e., how will a "controller" decide how to move among these four entities?). Modeling the environment is largely limited by technology, and recent advances in processing speeds and in computer graphics, displays and simulations (together called "hypermedia") make this a promising area for future design of AI tutors. Perhaps the most important task in designing the tutoring environment is to keep the student focused on the tutor's goal; the danger lies in making the environemnt too rich so that the student gets lost wandering through the hypermedia displays.

Perhaps the most difficult aspect to model is the student's cognitive processes. Despite the great strides cognitive science has made in the last decade in understanding the nature of learning, there is still much that is

not known. In order to build a model of the student's knowledge, the tutor needs information on a multitude of factors. For example, the tutor needs a method for deciding whether student errors are conceptual, strategic, procedural, or due simply to skills-deficiency (i.e. poor algebraic skills). The tutor also needs to know how much free reign to give the student, that is, how long should the tutor wait after an error is committed before interrupting the student; to interrupt too often interferes with learning, but to let the student go too far astray is also counterproductive. Tied to when and how often to interrupt the student is the student's intellectual and emotional profile--does the student lack motivation?, Is the student bright and motivated, and thus often bored with the pace of instruction or with the same tutorial strategy? Although not very helpful in answering these last questions, the research findings reviewed earlier are the basic ingredients needed to model the student's cognitive processes. The perfect design team for modeling the student's cogntive processes would include cognitive scientists, instructional designers, expert teachers and psychologists.

The fourth factor needed is a model of the tutoring strategies that the tutor will employ. Here one needs to decide how to structure the student-computer dialogue. Possible tutorial approaches include Socratic dialogue, analogical reasoning, immediate corrective feedback, worked-out examples, etc. Of course, the difficult part is deciding what particular combination of tutoring strategies work best for a particular student. This is something that the tutor must decide dynamically as it works with the student. If a particular tutoring strategy does appear to be ineffective, the tutor needs to switch to a different one. If one strategy was effective for a while but the student is no longer thriving under it, it may mean that the student is getting bored and losing motivation; the tutor needs to assess these situations and optimize on selecting the appropriate combination of tutoring strategies.

To date, many AI tutors of varying sophistication have been built in domains such as algebra (McArthur, Stasz & Hotta, 1987), geometry (Anderson, Boyle & Yost, 1985), electronic trouble-shooting (Brown, Burton & De Kleer, 1982; White & Frederiksen, 1986), medical diagnosis (Clancey, 1982, 1986), programming languages (Anderson & Reiser, 1986; Johnson & Soloway, 1984), military equipment maintenance (Towne & Munro, 1988), complex industrial processes (Woolf, Blegen, Jansen & Verloop, 1986), and various other domains. (For an overview of existing AI tutors, we refer the reader to Woolf, 1988). As one can surmise from the foregoing discussion, designing and building AI tutors is an expensive undertaking.

New Directions and Future Trends

We began this chapter by defining Information Processing as the link between Cognitive Science and Knowledge Engineering. The paradigm for IP has been language-based, by and large, as could be seen in the examples cited (e.g., how experts analyze their skills - Chi, et al, 1981; and how variable-reversal errors conform to surface-level characteristics of English

syntax - Mestre, 1988). New developments are beginning to explore other sensory processing approaches of IP, such as optical scanning, tactile responsiveness, auditory processes of echolocation, etc. While optical scanning is currently mechanical both in paradigm and in concept, visual IP and tactile IP are current research paradigms that hold considerable promise for future developments.

Clifton and Perris (1988), for example, is studying the roles of audition versus vision in the development of infants' guided reaching by attaching infrared light-emitting diodes to their fingers and video-recording reaching movements when the infant is in a darkened room. Reaching toward sound stimuli clearly cannot be visually-guided in such situations. The research will aid in understanding developmental sensory integration across projectile and ballistic reaching, vision, and auditory information processing.

Robotics is another application of Cognitive Science which is utilizing data from human skills of sensory information processing. Friedman & Cocking (1986) reviewed research on how blind persons were taught to identify and to locate in space complex forms, objects, figures, and faces (Back-y-Rita, 1980; 1981; 1982; 1983). The blind persons received visual information by controlling a television camera that delivered visual information to the skin through an array of vibratory stimulators or electrodes at the back, thigh, or abdomen. The authors reported that "the subjective perception of the obtained information was not on the skin; it was accurately located in the three-dimensional world in front of the camera" (Friedman & Cocking, 1986).

Exploration of the wide array of sensory experiences and their corresponding sensory systems will contribute to the Cognitive Science revolution that is now underway. The Information Processing framework is only beginning to define what constitutes "information." The associated information processing skills and strategies for achieving expert use of these skills in solving human factors-related problems are the challenge for Cognitive Science in Knowledge Engineering.

References

Anderson, J.R., Boyle, C. & Yost, G. (1985). The geometry tutor. Proceedings of the International Joint Conference on Artificial Intelligence, Los Angeles, CA. (pp. 1-7). Los Altos, CA: Morgan Kaufmann Publishers.

Anderson, J.R., & Reiser, B. (1986). The LISP tutor. Byte, 10, No. 4, 159-175.

Bach-y-Rita, P. (1980). Brain plasticity as a basis for therapeutic procedures. In P. Bach-y-Rita (Ed.), Recovery of function: Theoretical considerations for brain injury rehabilitation. Baltimore: University Park Press.

Bach-y-Rita, P. (1981). Brain plasticity as a basis of the development of rehabilitation procedures for hemiplegia. Scandinavian Journal of Rehabilitation Medicine, 13, 73-83.

Bach-y-Rita, P. (1982). The relationship between motor processes and cognition in tactile vision substitution. In W. Prinz & A. Sanders (Eds.), Cognition and motor processes. New York: Springer.

Bach-y-Rita, P. (1983). Tactile vision substitution: Past and future. International Journal of Neuroscience, 19, 29-36.

Beieiter, C. & Scarmadalia, M. (1978). From conversation to composition: The role of instruction in a developmental process. In R. Glaser (Ed.), Advances in instructional psychology. Hillsdale, NJ: Erlbaum.

Bower, G. H. Cognitive Psychology: An Introduction. In W. K. Estes (Ed.)Handbook of learning & cognitive processes, Vol. 1. Hillsdale, NJ: Erlbaum.

Bower, G. H., Black, J. B. & Turner, T. J. (1979). Scripts in memory for text. Cognitive Psychology. 11, 177-220.

Brown, A. L., Campione, J. C. & Day, J. D. (1981). Learning to learn: On training students to learn from texts. Educational Researcher, 10, 14-21.

Brown, J.S., Burton, R. & De Kleer, J. (1982). Pedagogical, natural language and knowledge engineering techniques in SOPHIE, I, II and III. (pp. 227-282). In D. Sleeman & J.S. Brown (Eds.), Intelligent Tutoring Systems. London: Academic Press.

Brown, J.S., Collins, A. & Duguid, P. (1989). Situated cognition and the culture of learning. Educational Researcher, 18, NO. 1, 32-42.

Chase, W.G. & Simon, H.A. (1973). Perception in chess. Cognitive Psychology, 4, 55-81.

Chi, M.T.H. & Glaser, R. (1981). The measurement of expertise: Analysis of

the development of knowledge and skills and a basis for assessing achievement. In E.L. Baker & E.S. Quellmalz (Eds.), Design, Analysis and Policy in Testing. Beverly Hills, CA: Sage Publications.

Chi, M.T.H., Feltovich, P. & Glaser, R. (1981). Categorization and representation of physics problems by experts and novices. Cognitive Science, 5, 121-152.

Clifton, R. & Perris, E. (1988) Infants' reaching in the dark reflects auditory development. Symposium presentation at Perceptual Development Conference. Jambaroo Valley, Wollongong, Australia.

Clancey, W.J. (1982). Tutoring rules for guiding a case method dialogue. In D.H. Fleeman & J.S. Brown (Eds.), Intelligent Tutoring Systems (pp. 201-225). London: Academic Press.

Clancey, W.J. (1986). From GUIDON to NEOMYCIN and HERACLES in twenty short lessons: ONR Final Report, 1979-1985. AI Magazine, 7, No. 3, 40-60.

Clement, J.J. (1982). Algebra word problem solutions: Thought processes underlying a common misconception. Journal for Research in Mathematics Education, 13, 16-30.

Cocking, R. R. (1983). Early concept formation. In E. K. Scholnick (Ed.), New trends in conceptual representation. Hillsdale, NJ: Erlbaum.

Cocking, R. R. & Chipman, S. F. (1988). Conceptual issues related to mathematics achievement of language minority children. In R. R. Cocking & J. P. Mestre (Eds.) Linguistic & Cultural Influences on Learning Mathematics. Hillsdale, NJ: Lawrence Erlbaum Associates.

de Groot, A.D. (1965). Thought and Choice in Chess. The Hague: Mouton.

Dufresne, R., Gerace, W., Hardiman, P. & Mestre, J. (in preparation). The effects of structured problem analysis on schema formation.

Egan, D.E. & Schwartz, B.J. (1979). Chunking in recall of symbolic drawings. Memory & Cognition, 7, 149-158.

Ehrlich, K. & Soloway, E. (1984). An empirical investigation of the tacit plan knowledge in programming. In J. Thomas & M.L. Schneider (Eds.), Human factors in computer systems. Norwood, NJ: Ablex.

Eylon, B. & Reif, F. (1984). Effects of knowledge organization on task performance. Cognition & Instruction, 1, 5-44.

Friedman, S. L. & Cocking, R. R. (1986). Instructional influences on cognition and on the brain. In S. L. Friedman, K. A. Klivington & R. W. Peterson (Eds.) The Brain, Cognition & Education. New York: Academic

Press.

Friedman, S. L., Scholnick, E. K. & Cocking, R. R. (1987) Blueprints for thinking: The role of planning in cognitive development. New York: Cambridge University Press.

Friedman, S. L., Scholnick, E. K. & Cocking, R. R. (1987). reflections on reflections: What planning is and how it develops. In S. L. Friedman, E. K. Scholnick & R. R. Cocking (Eds.), Blueprints for thinking: The role of planning in cognitive development. New York: Cambridge University Press.

Feuerstein, R., Rand, Y., Hoffman, M. B. & Miller, R. (1980). Instrumental enrichment: An Intervention Program for Cognitive Modifiability. Baltimore, MD: University Park Press.

Gardner, H. (1984). Frames of Mind. New York: Basic Books.

Gardner, H. (1986). Notes on cognitive development: Recent trends, new directions. In S. L. Friedman, K. A. Klivington & Peterson, R. W. (Eds.) the Brain, Cognition, & Education. New York: Academic Press.

Gerhard, C. (1975). Making Sense: Reading Comprehension Improved Through Categorizing. Newark, DE: International Reading Association.

Gibson, E. J. (1982). Perceiving as an adaptive process. Paper presented at the International Conference on Infant Studies, Austin, TX.

Gibson, J. J. (1966) The Senses Considered as Perceptual Systems. Boston: Houghton Mifflin.

Glaser, R. (1984). Education and thinking: The role of knowledge. American Psychologist, 39, 93-104.

Glaser, R. (1989). Expert knowledge and processes of thinking. Paper presented at the annual meeting of the American Association for the Advancement of Science, January 14-19, 1989, San Francisco, CA.

Hardiman, P.T., Dufresne, R. & Mestre, J.P. (in press). The relation between problem categorization and problem solving among experts and novices. Memory & Cognition.

Heller, J. & Reif, F. (1984). Prescribing effective human problem solving processes: Problem description in physics. Cognition & Instruction, 1, 177-216.

Johnson, W.L. & Soloway, E.M. (1984). Intention-based diagnosis of programming errors. Proceedings of the National Conference on Artificial Intelligence, Austin, TX (pp. 162-168). Los Altos, CA: Morgan Kaufmann Publishers.

Larkin, J.H. (1979). Information processing models in science instruction. In J. Lochhead & J. Clement (Eds.), Cognitive Process Instruction.

Hillsdale, NJ: Lawrence Erlbaum Associates.

Larkin, J., Heller, L. & Greeno, J. (1980) Instructional implications of research on problem solving. In W. J. McKeachie (Ed.), Cognition, college teaching, and student learning. San Francisco: Jossey-Bass.

Larkin, J., McDermott, J., Simon, H. & Simon, D. (1980a). Expert and novice performance in solving physics problems. Science, 208, 1335-1342.

Larkin, J., McDermott, J., Simon, H. & Simon, D. (1980b). Models of competence in solving physics problems. Cognitive Science, 4, 317-345.

McArthur, D., Stasz, C. & Hotta, J. (1987). Learning problem solving skills in algebra. Journal of Educational Technology Systems, 15, No. 3, 304-324.

Meichenbaum, D. (1977). Cognitive Behavior & Modification: An Integrative Approach. New York: Plenum.

Mestre, J.P. (1988). The role of language comprehension in mathematics and problem solving. In R. Cocking & J. Mestre (Eds.), Linguistic and Cultural Influences in Learning Mathematics (pp. 201-220). Hillsdale, NJ: Lawrence Erlbaum Associates.

Mestre, J., Dufresne, R., Gerace, W. & Hardiman, P. (1988). Hierarchical problem solving as a means of promoting expertise. Proceedings of the Tenth Annual Conference of the Cognitive Science Society (pp. 312-318). Hillsdale, NJ: Lawrence Erlbaum Associates.

Nelson, K. (1978). How young children organize knowledge of the world in and out of language. In R. Siegler (Ed.) Children's Thinking: What Develops? Hillsdale, NJ: Erlbaum.

Nelson, K. (1982). Social cognition in a script framework. In J. H. Flavell & L. Ross (Eds.) Social cognitive development. New York: Cambridge University Press.

Newell, A. & Simon, H. (1972). Human Problem Solving. Englewood Cliffs, NJ: Prentice Hall.

Nickerson, R.S., Perkins, D. N. & Smith, E. E. (1980). Teaching Thinking. Cambridge, MA: Bolt, Beranek & Newman.

Owen, E. & Sweller, J. (1985). What do students learn while solving mathematics problems? Journal of Educational Psychology, 77, 272-284.

Perkins, D. & Salomon, G. (1989). Are cognitive skills context-bound? Educational Researcher, 18, NO. 1, 16-25.

Rabinowitz, M. & Glaser, R. (1985). Cognitive structure and process in highly competent performance. In F.D. Horowitz & M. O'Brian (Eds.), The

Gifted and Talented: Developmental Perspectives (pp. 75-98). Washington, D.C.: American Psychological Association.

Rogoff, B. (1981). Schooling and the development of cognitive skills. In H. C. Triandis & A. Heron (Eds.) Handbook of cross-cultural psychology (Vol. 4). Rockleigh, NJ: Allyn & Bacon.

Rumelhart, D. E. & Norman, D. A. (1975). The active structural network. In D. A. Normal & D. E. Rumelhart (Eds.) Explorations in cognition. San Francisco: W. H. Freeman.

Schoenfeld, A. (1985). Mathematical Problem Solving. New York, NY: Academic Press.

Schoenfeld, A.H. & Herrmann, D.J. (1982). Problem perception and knowledge structure in expert and novice mathematical problem solvers. Journal of Experimental Psychology: Learning, Memory and Cognition, 8, 484-494.

Shank, R. C. (1975). The structure of episodes in memory. In D. G. Bobrow & A. Collins (Eds.) Representation and understanding. New York: Academic Press.

Shank, R. C. & Abelson, R. (1977). Scripts, plans, goals, and understanding. Hillsdale, NJ: Erlbaum.

Sigel, I. E. & Cocking R. R. (1977). Cognitive Development from Childhood to Adolescence: A Constructivist Perspective. New York: Holt, Rinehart & Winston.

Simon, D.P. & Simon, H.A. (1978). Individual differences in solving physics problems. In R.S. Sigler (Ed.), Children's Thinking: What Develops? Hillsdale, NJ: Lawrence Erlbaum Associates.

Sweller, J., Mawer, R. & Ward, M. (1983). Development of expertise in mathematical problem solving. Journal of Experimental Psychology: General. 112, 639-661.

Sweller, J. (1988). Cognitive load during problem solving: Effects on learning. Cognitive Science, 12, 157-285.

Towne, E.,D. & Munro, A. (1988). The intelligent maintenance training system. In J. Psotka, L. Massey & S. Mutter (Eds.), Intelligent Tutoring Systems: Lessons Learned (pp. 479-530). Hillsdale, NJ: Lawrence Erlbaum Associates.

Touger, J., Dufresne, R., Gerace, W. & Mestre, J. (1987). Hierarchical organization of knowledge and coherent explanation in the domain of elementary mechanics. In J.D. Novak (Ed.), Proceedings of the Second International Seminar on Misconceptions and Educational Strategies in Science and Mathematics, Vol III (pp. 517-530). Ithaca, NY: Department

of Education, Cornell University.

Vygotsky, L. S. (1978). Mind & Society. Cambridge, MA: Harvard University Press.

White, B.Y. & Frederiksen, J. (1986). Intelligent tutoring systems based upon qualitative model evolution. Proceedings of the National Conference on Artificial Intelligence, Philadelphia, PA (pp. 313-319). Los Altos, CA: Morgan Kaufmann Publishers.

Woolf, B.P. (1988). Intelligent tutoring systems: A survey. In H.E. Shrobe (Ed.), Exploring Artificial Intelligence: Survey Talks from the National Conferences on Artificial Intelligence (pp. 1-43). San Mateo, CA: Morgan Kaufmann Publishers.

Woolf, B.P., Blegen, D., Jansen, J. & Verloop, A. (1986). Teaching a complex industrial process. Proceedings of the National Conference on Artificial Intelligence, Philadelphia, PA (pp. 722-727). Los Altos, CA:Morgan Kaufmann Publishers.